

CLAIMS

1. A system for failover comprising:

at least one client selectively connectable to at least one of a primary server and at least one backup server via a network connection;

5 said primary server and said at least one backup server being connected to each other in a normal state during which said primary server processes messages from said client and maintains a transaction log respective to each said message in random access memory, and during which said backup server is operable to maintain a mirror image of said transaction log, said transaction logs being periodically flushed to a non-volatile storage device respective to each said server;

10 said client being only connected to said primary server in a primary-only state when said backup server is unavailable during which said primary server processes messages from said client and maintains a transaction log respective to each said message in random access memory, and during which said primary server transaction log is flushed to said primary server non-volatile storage device after each said message is processed; and,

15 said client being only connected to said backup server in a backup-only state when said primary server is unavailable during which said backup server initially recovers a last-known state of said primary server and then processes messages from said client and maintains said backup server transaction log respective to each said message in random access memory, and wherein said transaction log is flushed to said backup server non-volatile storage device immediately after each said message is processed.

2. A system for failover comprising:

at least one client, a primary server and at least one backup server interconnected via at least one network connection;

25 said primary server and said at least one backup server each operable to:

execute an application for processing messages received from said at least one client;

maintain a transaction log in random access memory;

execute a replication agent;

execute a checkpoint agent for flushing said transaction log to a non-volatile storage device respective to each said server;

5 said system having:

a normal state during which only said primary server application processes said messages and writes results thereof to said primary server transaction log and during which said replication agents are operable to mirror said primary server transaction log in said backup server transaction log, each of said transaction logs being flushed according to a predefined criteria;

10

a primary-only state during which only said primary server application processes said messages and writes results thereof to said primary server transaction log and during which transaction log is simultaneously flushed to said primary server non-volatile storage device;

15

a backup-only state during which only said backup server application processes said messages and writes results thereof to said backup server transaction log and during which backup server transaction log is simultaneously flushed to said backup server non-volatile storage device.

3. A system for failover comprising:

20

at least one client connectable to a primary server coupled to a backup server in a normal state during which said primary server processes messages from said client and maintains a transaction log respective to each said message in random access memory, and during which said backup server is operable to maintain a mirror image of said transaction log, said transaction logs being periodically flushed to a non-volatile storage device respective to each said server;

25

5 said at least one client only connectable to said primary server when said backup server is unavailable during which said primary server processes messages from said client and maintains a transaction log respective to each said message in random access memory, and during which said backup server transaction log is flushed to said primary server non-volatile storage device after each said message is processed; and,

10 said at least one client only connectable to said backup server when said primary server is unavailable during which said backup server initially recovers a last-known state of said primary server and then processes messages from said client and maintains said backup server transaction log respective to each said message in random access memory, and wherein said transaction log is flushed to said backup server non-volatile storage device immediately after each said message is processed.

4. A system for failover comprising:

15 at least one client connectable to a primary server coupled to a backup server when both of said servers are available and during which said primary server processes messages from said client and maintains a transaction log respective to each said message in random access memory, and during which said backup server is operable to maintain a mirror image of said transaction log, said transaction logs being periodically flushed to a non-volatile storage device respective to each said server;

20 said at least one client being connectable to an available one of said primary server and said backup server when one of said servers is unavailable, during which said available server processes messages from said client and maintains a transaction log respective to each said message in random access memory, and during which said available server transaction log is flushed to said non-volatile storage device respective thereto after each said message is processed.

25 5. A system for failover comprising at least one client connectable to a plurality of interconnectable servers, each of said servers operable to process messages from said client and maintain at least transaction record respective to each said message in a transaction log stored in volatile storage, and further operable to flush said transaction log to a non-volatile storage device

respective to said server after each said message is processed, wherein when more than one server is available then only one of said servers processes said messages and a remainder of said servers maintain a mirror image of said transaction log, said transaction logs being flushed on a periodic basis, and wherein when only one of said servers is available then said transaction log is flushed more frequently than said periodic basis.

6. The system according to claim 5 wherein said periodic basis is an interval based on one of the criteria selected from the group consisting of a fixed time interval; a predetermined number of said transaction records having been stored in said volatile storage; and a time when said volatile storage of a respective one of said servers needs to be freed for other purposes.

10 7. The system according to claim 5 wherein said servers are connectable to said clients via first network and wherein said servers are interconnectable with each other via a second network.

8. The system according to claim 7 wherein said second network comprises a main link and a failsafe link.

15 9. The system according to claim 7 wherein said first network is selected from the group consisting of the Internet, a local area network and a wide area network.

10. The system according to claim 5 wherein said client is selected from the group consisting of a personal digital assistant, a cellular telephone, a laptop computer and an email paging device.

20 11. The system according to claim 5 wherein said servers include a primary server and a backup server

12. The system according to claim 1 wherein said primary server and said backup server are implemented on two substantially identical computing environments.

13. The system according to claim 13 comprising at least one additional backup server.

25 14. The system according to claim 11 wherein said primary server and said backup server are implemented on a single computing environment that is logically partitioned into said primary server and said backup server.

15. The system according to claim 11 wherein said servers include a plurality of software components comprising at least one application process that operates in a hot state to process said messages and otherwise operates in a warm state.

16. The system according to claim 15 wherein when said primary server is available then
5 said primary server application process operates in said hot state, and when said primary server is unavailable then said backup servers bring its application process from said warm state to said hot state and is thereafter operable to process said messages from said clients.

17. The system according to claim 15 wherein said software components include a failover agent executing on each of said servers and communicating with each to verify whether each
10 other said server is available.

18. The system according to claim 17 wherein when said primary server failover agent fails to communicate with said backup server failover agent then said primary server begins operating in a primary-only mode during which said primary server transaction log is flushed after each message.

15 19. The system according to claim 17 wherein when said backup server failover agent fails to communicate with said primary server failover agent then said backup server begins operating in a backup-only mode during which said backup server transaction log is flushed after each message.

20 20. The system according to claim 15 wherein said software components include a replication agent executing on each of said servers and communicating with each other to maintain said mirror image when both of said servers are available.

21. The system according to claim 15 wherein said software components include a checkpoint agent executing on each of said servers and communicating with its respective transaction log, said checkpoint agent for flushing its respective transaction log to non-volatile
25 storage.

22. The system according to claim 5 wherein said servers are operable to determine whether a received one of said messages is a duplicate and to return a result of a stored record representing a previous processing of said message in lieu of processing said message.

23. The system according to claim 5 wherein said message includes an identifier of said client, a message identifier, and a message body.

24. The system according to claim 23 wherein said message body includes a request to purchase or sell a financial instrument.

25. The system according to claim 5 wherein said message includes a request that is processed as at least one discrete transaction, each of said transactions generating at least one transaction log record.

26. The system according to claim 11 wherein when said backup server is processing requests, said backup server is operable to assume the role of said primary server when said primary server becomes available and at which point said primary server assumes the role of said backup server.

27. A method for operating a system for failover comprising the steps of:
receiving, at a primary server, a request from a client;
processing, at said primary server, said request as at least one transaction;
saving, in volatile memory of said primary server, a transaction record respective to said at least one transaction;
generating, at said primary server, a request to commit said transaction;
mirroring said transaction record in volatile memory of said backup server; and,
acknowledging, in said primary server, said commit request if said mirroring is successfully confirmed.

28. The method of claim 27 further comprising the steps of:

periodically flushing said primary server volatile memory to a non-volatile memory respective to said primary server.

29. The method of claim 27 further comprising the steps of:

5 periodically flushing said backup server volatile memory to a non-volatile memory respective to said backup server.

30. The method of claim 27 further comprising the additional steps of repeating the foregoing steps for additional received messages.

31. The method of claim 30 further comprising the steps of:

10 immediately flushing said primary server volatile memory to non-volatile memory if said mirroring is unsuccessfully confirmed; and,

thereafter ceasing further performance of said mirroring step; and,

thereafter flushing each said transaction record to non-volatile memory after said saving step.

15 32. The method of claim 27 further comprising, after said receiving step, the step of determining whether said request is a duplicate request, and if said request is a duplicate request, then returning, to said client, a previously stored result of processing said request and then omitting the remaining steps of said method prior to repeating said method.

33. A method of operating a backup server in a system for failover comprising the steps of:

20 when said backup server determines a primary server connected thereto is available, performing the steps of:

mirroring transaction records generated by said primary server in volatile memory of said backup server; and,

periodically flushing said volatile memory to non-volatile memory connected to said backup server;

when said backup server determines said primary server is unavailable, performing the steps of:

receiving a request from a client;

processing said request as at least one transaction;

5 saving in volatile memory a transaction record respective to said at least one transaction; and,

flushing said at least one transaction record to non-volatile memory.

34. The method of claim 33 further comprising the steps of, when said backup server initially determines said primary server is no longer available;

10 flushing said backup server volatile memory to non-volatile memory;

applying any outstanding transactions to at least a table and a queue associated with said transaction records stored on said non-volatile memory; and,

sending a signal to said clients indicating said backup server is ready to process requests.

35. A backup server in a system for failover comprising:

15 means for, when said backup server determines a primary server connected thereto is available, performing the steps of:

mirroring transaction records generated by said primary server in volatile memory of said backup server; and,

20 periodically flushing said volatile memory to non-volatile memory connected to said backup server;

means for, when said backup server determines said primary server is unavailable, performing the steps of:

receiving a request from a client;

processing said request as at least one transaction;

saving in volatile memory a transaction record respective to said at least one transaction;
and,

flushing said at least one transaction record to non-volatile memory.

5 36. The server of claim 35 further comprising means for performing the steps of, when said backup server initially determines said primary server is no longer available;

flushing said backup server volatile memory to non-volatile memory;

applying any outstanding transactions to at least a table and a queue associated with said transaction records stored on said non-volatile memory; and,

10 sending a signal to said clients indicating said backup server is ready to process requests.

37. A primary server in a system for failover comprising:

means for receiving a request from a client;

means for processing said request as at least one transaction;

means for saving in volatile memory a transaction record respective to said at least one transaction;

means for generating a request to commit said transaction;

means for sending a request to mirror said transaction record in volatile memory of a backup server; and,

acknowledging said commit request if said mirroring request is successfully confirmed.

20 38. The server of claim 37 further comprising:

means for periodically flushing said primary server volatile memory to a non-volatile memory respective to said primary server.